

Joint Detection and Identification of an Unobservable Change in the Distribution of a Random Sequence

Savas Dayanik and Christian Goulding

Dept. of Operations Research and Financial Engineering

Princeton University, Princeton, NJ 08544

Email: {sdayanik, cgouldin}@princeton.edu

H. Vincent Poor

School of Engineering and Applied Science

Princeton University, Princeton, NJ 08544

Email: poor@princeton.edu

Abstract—This paper examines the joint problem of detection and identification of a sudden and unobservable change in the probability distribution function (pdf) of a sequence of independent and identically distributed (i.i.d.) random variables to one of finitely many alternative pdf's. The objective is quick detection of the change and accurate inference of the ensuing pdf. Following a Bayesian approach, a new sequential decision strategy for this problem is revealed and is proven optimal. Geometrical properties of this strategy are demonstrated via numerical examples.

I. INTRODUCTION

Consider a sequence of i.i.d. random variables X_1, X_2, \dots , taking values in some measurable space (E, \mathcal{E}) . The common probability distribution of the X 's is initially some known probability measure \mathbb{P}_0 on (E, \mathcal{E}) , and then, at some *unobservable* disorder time θ , the common probability distribution changes suddenly to another probability measure \mathbb{P}_μ for some *unobservable* index $\mu \in \mathcal{M} \triangleq \{1, \dots, M\}$. The objective is to detect the change as quickly as possible, and, at the same time, to identify the new probability distribution as accurately as possible, so that the most suitable actions can be taken with the least delay.

This problem can be viewed as the fusion of two fundamental areas of sequential analysis: change detection and multi-hypothesis testing. In traditional change detection problems, there is only one change distribution, \mathbb{P}_1 ; therefore, the focus is exclusively on detecting the change time. Whereas, in traditional sequential multi-hypothesis testing problems, there is no change time to consider. Instead, every observation has common distribution \mathbb{P}_μ for some unknown μ , and the focus is exclusively on the inference of μ . Both of these subproblems have been studied extensively. For recent reviews of these areas, we refer the reader to [1] and [2] and the references therein.

However, the joint problem involves key trade-off decisions not taken into account by separately applying techniques for these subproblems. While raising an alarm as soon as the change occurs is advantageous for the change detection task, it is undesirable for the identification task because waiting longer provides more observations for inferring the change distribution. Likewise, the unknown change time complicates the identification task, and, as a result, adaptation of existing sequential multi-hypothesis testing algorithms is problematic.

The research of Savas Dayanik was supported by the Air Force Office of Scientific Research, under grant AFOSR-FA9550-06-1-0496. The research of H. Vincent Poor was supported in part by the U.S. Army Pantheon Project.

Decision strategies for the joint problem have a wide array of applications, such as fault detection and isolation in industrial processes, target detection and identification in national defense, pattern recognition and machine learning, radar and sonar signal processing, seismology, speech and image processing, biomedical signal processing, finance, and insurance. However, the theory has not been broadly developed. Nikiforov [3] provides the first results for this problem, showing asymptotic optimality for a certain non-Bayesian approach, and Lai [4] generalizes these results through the development of information-theoretic bounds and the application of likelihood methods. In this paper, we follow a Bayesian approach to reveal a new optimal strategy for this problem and we describe an accurate numerical scheme for its implementation.

In Sec. II we formulate precisely the problem in a Bayesian framework, and in Sec. III we show that it can be reduced to an optimal stopping of a Markov process whose state space is the standard probability simplex. In addition, we establish a simple recursive formula that captures the dynamics of the process and yields a sufficient statistic fit for online tracking.

In Sec. IV we use optimal stopping theory to substantiate the optimality equation for the value function of the optimal stopping problem. Moreover, we prove that this value function is bounded, concave, and continuous on the standard probability simplex and that the optimal stopping region consists of M non-empty, convex, closed, and bounded subsets. Also, we consider a truncated version of the problem that allows at most N observations from the sequence of random measurements. We establish an explicit bound (inversely proportional to N) for the approximation error associated with this truncated problem.

In Sec. V we show that the separate problems of change detection and sequential multi-hypothesis testing are solved as special cases of the overall joint solution. We illustrate some geometrical properties of the optimal method and demonstrate its implementation by numerical examples for the special cases $M = 2$ and $M = 3$. Specifically, we show instances in which the M convex subsets comprising the optimal stopping region are connected and instances in which they are not. Likewise, we show that the continuation region (i.e., the complement of the stopping region) need not be connected. We refer the reader to [5] for complete proofs of the results.

II. PROBLEM STATEMENT

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space hosting random variables $\theta : \Omega \mapsto \{0, 1, \dots\}$ and $\mu : \Omega \mapsto \mathcal{M} \triangleq \{1, \dots, M\}$ and a process $X = (X_n)_{n \geq 1}$ taking values in some measurable

space (E, \mathcal{E}) . Suppose that for every $t \geq 1, i \in \mathcal{M}, n \geq 1$, and $(E_k)_{k=1}^n \subseteq \mathcal{E}$ we have

$$\begin{aligned} & \mathbb{P}\{\theta = t, \mu = i, X_1 \in E_1, \dots, X_n \in E_n\} \\ &= (1 - p_0)(1 - p)^{t-1} p \nu_i \prod_{k=1}^{(t-1) \wedge n} \mathbb{P}_0(E_k) \prod_{\ell=t \vee 1}^n \mathbb{P}_i(E_\ell) \end{aligned}$$

for some given probability measures $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_M$ on (E, \mathcal{E}) , known constants $p_0 \in [0, 1]$, $p \in (0, 1)$, and $\nu_i > 0, i \in \mathcal{M}$ such that $\nu_1 + \dots + \nu_M = 1$, where $x \wedge y \triangleq \min\{x, y\}$ and $x \vee y \triangleq \max\{x, y\}$. Namely, θ is independent of μ ; it has a zero-modified geometric distribution with parameters p_0 and p in the terminology of [6, Sec. 3.6], which reduces to the standard geometric distribution when $p_0 = 0$.

Conditionally on θ and μ , the random variables $X_n, n \geq 1$ are independent; $X_1, \dots, X_{\theta-1}$ and $X_\theta, X_{\theta+1}, \dots$ are identically distributed with common distributions \mathbb{P}_0 and \mathbb{P}_μ , respectively. The probability measures $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_M$ always admit densities with respect to some sigma-finite measure m on (E, \mathcal{E}) ; for example, we can take $m = \mathbb{P}_0 + \mathbb{P}_1 + \dots + \mathbb{P}_M$. So, we fix m and denote the corresponding densities by f_0, f_1, \dots, f_M , respectively.

Suppose now that we observe sequentially the random variables $X_n, n \geq 1$. Their common pdf f_0 changes at stage θ to some other pdf $f_\mu, \mu \in \mathcal{M}$. Our objective is to detect the change time θ as quickly as possible *and* to identify the change index μ as accurately as possible. More precisely, given costs associated with detection delay, false alarm, and false identification of the change index, we seek a strategy that minimizes the expected total change detection *and* identification cost.

Let $\mathbb{F} = (\mathcal{F}_n)_{n \geq 0}$ denote the natural filtration of the observation process X , where

$$\mathcal{F}_0 = \{\emptyset, \Omega\} \quad \text{and} \quad \mathcal{F}_n = \sigma(X_1, \dots, X_n), \quad n \geq 1.$$

A strategy $\delta = (\tau, d)$ is a pair consisting of a *stopping time* τ of the filtration \mathbb{F} and a *terminal decision rule* $d : \Omega \mapsto \mathcal{M}$ measurable with respect to the history $\mathcal{F}_\tau = \sigma(X_{n \wedge \tau}; n \geq 1)$ of observation process X through stage τ . Applying a strategy $\delta = (\tau, d)$ consists of announcing at the end of stage τ that the common pdf has changed from f_0 to f_d at or before stage τ . Let

$$\Delta \triangleq \{(\tau, d) \mid \tau \in \mathbb{F}, \text{ and } d \in \mathcal{F}_\tau \text{ is an } \mathcal{M}\text{-valued r. v.}\}$$

denote the collection of all such sequential decision strategies.

For every strategy $\delta = (\tau, d) \in \Delta$, we define a *Bayes risk function*

$$R(\delta) = c \mathbb{E}[(\tau - \theta)^+] + \mathbb{E}[a_{0d} \mathbf{1}_{\{\tau < \theta\}} + a_{\mu d} \mathbf{1}_{\{\theta \leq \tau < \infty\}}] \quad (1)$$

as the expected diagnosis cost: the sum of the expected detection delay cost and the expected terminal decision cost upon alarm, where $c > 0$ and $a_{ij} \geq 0, i \in \{0\} \cup \mathcal{M}, j \in \mathcal{M}$ are known constants satisfying $a_{ii} = 0, i \in \mathcal{M}$ (i.e., no cost for a correct terminal decision), and $(x)^+ \triangleq \max\{x, 0\}$.

The problem is to find a sequential decision strategy $\delta = (\tau, d) \in \Delta$ (if it exists) with the *minimum Bayes risk*

$$R^* \triangleq \inf_{\delta \in \Delta} R(\delta). \quad (2)$$

III. POSTERIOR ANALYSIS AND FORMULATION AS AN OPTIMAL STOPPING PROBLEM

In this section we show that the Bayes risk function in (1) can be written as the expected value of the running and terminal costs driven by a certain Markov process. We use this fact to recast the minimum Bayes risk in (2) as a Markov optimal stopping problem.

Let us introduce the posterior probability processes

$$\Pi_n^{(0)} \triangleq \mathbb{P}\{\theta > n \mid \mathcal{F}_n\} \quad \text{and} \quad \Pi_n^{(i)} \triangleq \mathbb{P}\{\theta \leq n, \mu = i \mid \mathcal{F}_n\}$$

for $i \in \mathcal{M}, n \geq 0$. Having observed the first n observations, $\Pi_n^{(0)}$ is the posterior probability that the change *has not* yet occurred at or before stage n , while $\Pi_n^{(i)}$ is the posterior joint probability that the change *has* occurred by stage n and that the hypothesis $\mu = i$ is correct. The connection of these posterior probabilities to the loss structure for our problem is established in the next proposition.

Proposition 1. *For every sequential decision strategy $\delta \in \Delta$, the Bayes risk function (1) can be expressed in terms of the process $\Pi \triangleq \{\Pi_n = (\Pi_n^{(0)}, \dots, \Pi_n^{(M)})\}_{n \geq 0}$ as*

$$R(\delta) = \mathbb{E} \left[\sum_{n=0}^{\tau-1} c (1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} \sum_{j=1}^M \mathbf{1}_{\{d=j\}} \sum_{i=0}^M a_{ij} \Pi_\tau^{(i)} \right].$$

While our original formulation of the Bayes risk function (1) was in terms of the values of the unobservable random variables θ and μ , Proposition 1 gives us an equivalent version of the Bayes risk function in terms of the posterior distributions for θ and μ . This is particularly effective in light of Proposition 2, which we state with the aid of some additional notation that is referred to throughout the paper. Let

$$S^M \triangleq \left\{ \pi = (\pi_0, \pi_1, \dots, \pi_M) \in [0, 1]^{M+1} \mid \sum_{i=0}^M \pi_i = 1 \right\}$$

denote the standard M -dimensional probability simplex. Define the mappings $D_i : S^M \times E \mapsto [0, 1], i \in \mathcal{M}$ and $D : S^M \times E \mapsto [0, 1]$ by

$$D_i(\pi, x) \triangleq \begin{cases} (1 - p)\pi_0 f_0(x), & i = 0 \\ (\pi_i + \pi_0 p \nu_i) f_i(x), & i \in \mathcal{M} \end{cases}$$

and $D(\pi, x) \triangleq \sum_{i=0}^M D_i(\pi, x)$, and the operator \mathbb{T} on the collection of bounded functions $f : S^M \mapsto \mathbb{R}$ by

$$(\mathbb{T}f)(\pi) \triangleq \int_E m(dx) D(\pi, x) f\left(\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)}\right) \quad (3)$$

for every $\pi \in S^M$.

Proposition 2. (a) *The process $\Pi^{(0)} \triangleq \{\Pi_n^{(0)}, \mathcal{F}_n\}_{n \geq 0}$ is a supermartingale, and $\mathbb{E} \Pi_n^{(0)} \leq (1 - p)^n$ for every $n \geq 0$.*

(b) *The process $\Pi^{(i)} \triangleq \{\Pi_n^{(i)}, \mathcal{F}_n\}_{n \geq 0}$ is a submartingale for every $i \in \mathcal{M}$.*

(c) *The process $\Pi = \{\Pi_n^{(0)}, \dots, \Pi_n^{(M)}\}_{n \geq 0}$ is a Markov process, and*

$$\Pi_{n+1}^{(i)} = \frac{D_i(\Pi_n, X_{n+1})}{D(\Pi_n, X_{n+1})}, \quad i \in \{0\} \cup \mathcal{M}, \quad n \geq 0, \quad (4)$$

with initial state $\Pi_0^{(0)} = 1 - p_0$ and $\Pi_0^{(i)} = p_0 \nu_i$, $i \in \mathcal{M}$. Moreover, for every bounded function $f : S^M \mapsto \mathbb{R}$ and $n \geq 0$, we have $\mathbb{E}[f(\Pi_{n+1})|\Pi_n] = (\mathbb{T}f)(\Pi_n)$.

Remark 3. Since Π is uniformly bounded, the limit $\lim_{n \rightarrow \infty} \Pi_n$ exists by the martingale convergence theorem. Moreover, $\lim_{n \rightarrow \infty} \Pi_n^{(0)} = 0$ a.s. by Proposition 2(a) since $p \in (0, 1)$.

Now, let the functions h, h_1, \dots, h_M from S^M into \mathbb{R}_+ be defined by

$$h(\pi) \triangleq \min_{j \in \mathcal{M}} h_j(\pi) \quad \text{and} \quad h_j(\pi) \triangleq \sum_{i=0}^M \pi_i a_{ij}, \quad j \in \mathcal{M},$$

respectively. Then, we note that for every $\delta = (\tau, d) \in \Delta$, we have

$$\begin{aligned} R(\tau, d) &= \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} \sum_{j=1}^M \mathbf{1}_{\{d=j\}} h_j(\Pi_\tau) \right] \\ &\geq \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} h(\Pi_\tau) \right] = R(\tau, \tilde{d}) \end{aligned}$$

where we define on the event $\{\tau < \infty\}$ the terminal decision rule \tilde{d} to be any index satisfying $h_{\tilde{d}}(\Pi_\tau) = h(\Pi_\tau)$. In other words, an optimal terminal decision depends only upon the value of the Π process at the stage in which we stop. Note also that the functions h and h_1, \dots, h_M are bounded on S^M . Therefore, we have the following:

Lemma 4. The minimum Bayes risk (2) reduces to the following optimal stopping of the Markov process Π :

$$\begin{aligned} R^* &= \inf_{(\tau, d) \in \Delta} R(\tau, d) = \inf_{(\tau, \tilde{d}) \in \Delta} R(\tau, \tilde{d}) \\ &= \inf_{\tau \in \mathbb{F}} \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} h(\Pi_\tau) \right]. \end{aligned}$$

We simplify this formulation further by showing that it is enough to take the infimum over

$$C \triangleq \{\tau \in \mathbb{F} \mid \tau < \infty \text{ a.s. and } \mathbb{E}Y_\tau^- < \infty\}, \quad (5)$$

where we define

$$-Y_n \triangleq \sum_{k=0}^{n-1} c(1 - \Pi_k^{(0)}) + h(\Pi_n), \quad n \geq 0$$

as the minimum *partial risk* obtained by making the best terminal decision on $\{\tau = n\}$. Since $h(\cdot)$ is bounded on S^M , the process $\{Y_n, \mathcal{F}_n; n \geq 0\}$ consists of integrable random variables. So the expectation $\mathbb{E}Y_\tau$ exists for every $\tau \in \mathbb{F}$, and our problem becomes

$$-R^* = \sup_{\tau \in \mathbb{F}} \mathbb{E}Y_\tau. \quad (6)$$

Observe that $\mathbb{E}\tau < \infty$ for every $\tau \in C$ because $\infty > (1/c)\mathbb{E}Y_\tau^- \geq \mathbb{E}(\tau - \theta)^+ \geq \mathbb{E}(\tau - \theta) \geq \mathbb{E}\tau - \mathbb{E}\theta \geq \mathbb{E}\tau - (1/p)$. In fact, we have $\mathbb{E}Y_\tau > -\infty \Leftrightarrow \mathbb{E}Y_\tau^- < \infty \Leftrightarrow \mathbb{E}\tau < \infty$ for every $\tau \in \mathbb{F}$. Since $\sup_{\tau \in \mathbb{F}} \mathbb{E}Y_\tau \geq \mathbb{E}Y_0 > -h(\Pi_0) > -\infty$, it

is enough to consider $\tau \in \mathbb{F}$ such that $\mathbb{E}\tau < \infty$. Namely, (6) reduces to

$$-R^* = \sup_{\tau \in C} \mathbb{E}Y_\tau. \quad (7)$$

IV. SOLUTION VIA OPTIMAL STOPPING THEORY

In this section we derive an optimal solution for the problem in (2) by building on the formulation of (7) via the tools of optimal stopping theory, which are detailed in [7].

A. The optimality equation.

We begin by applying the method of truncation with a view of passing to the limit to arrive at the final result. Define for every pair of integers n, N satisfying $0 \leq n \leq N$ the sub-collections

$$C_n \triangleq \{\tau \vee n \mid \tau \in C\} \quad \text{and} \quad C_n^N \triangleq \{\tau \wedge N \mid \tau \in C_n\}$$

of stopping times in C of (5) and the families of (truncated) optimal stopping problems

$$-V_n \triangleq \sup_{\tau \in C_n} \mathbb{E}Y_\tau \quad \text{and} \quad -V_n^N \triangleq \sup_{\tau \in C_n^N} \mathbb{E}Y_\tau \quad (8)$$

corresponding to $(C_n)_{n \geq 0}$ and $(C_n^N)_{0 \leq n \leq N}$, respectively. Note that $C \equiv C_0$ and $R^* \equiv V_0$.

To investigate these optimal stopping problems, we introduce versions of the *Snell envelope* of $(Y_n)_{n \geq 0}$ (i.e., the smallest regular supermartingale dominating $(Y_n)_{n \geq 0}$) corresponding to $(C_n)_{n \geq 0}$ and $(C_n^N)_{0 \leq n \leq N}$, respectively, defined by

$$\gamma_n \triangleq \operatorname{ess\,sup}_{\tau \in C_n} \mathbb{E}[Y_\tau \mid \mathcal{F}_n] \quad \text{and} \quad \gamma_n^N \triangleq \operatorname{ess\,sup}_{\tau \in C_n^N} \mathbb{E}[Y_\tau \mid \mathcal{F}_n].$$

Then through the following series of lemmas we point out several useful properties of these Snell envelopes. Finally, we extend these results to an arbitrary initial state vector and establish the optimality equation. Note that each of the ensuing (in)equalities between random variables are in the \mathbb{P} -almost sure sense.

First, these Snell envelopes provide the following alternative expressions for the optimal stopping problems introduced in (8) above.

Lemma 5. For every $N \geq 0$ and $0 \leq n \leq N$, we have $-V_n = \mathbb{E}\gamma_n$ and $-V_n^N = \mathbb{E}\gamma_n^N$.

Second, we have the following backward-induction equations.

Lemma 6. We have $\gamma_n = \max\{Y_n, \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n]\}$ for every $n \geq 0$. For every $N \geq 1$ and $0 \leq n \leq N-1$, we have $\gamma_N^N = Y_N$ and $\gamma_n^N = \max\{Y_n, \mathbb{E}[\gamma_{n+1}^N \mid \mathcal{F}_n]\}$.

We also have that these versions of the Snell envelopes coincide in the limit as $N \rightarrow \infty$. That is,

Lemma 7. For every $n \geq 0$, we have $\gamma_n = \lim_{N \rightarrow \infty} \gamma_n^N$.

Next, recall from (3) and Proposition 2(c) the operator \mathbb{T} and let us introduce the operator \mathbb{M} on the collection of bounded functions $f : S^M \mapsto \mathbb{R}_+$ defined by

$$(\mathbb{M}f)(\pi) \triangleq \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}f)(\pi)\}, \quad \pi \in S^M.$$

Observe that $0 \leq \mathbb{M}f \leq h$. That is, $\pi \mapsto (\mathbb{M}f)(\pi)$ is a nonnegative bounded function. Therefore, $\mathbb{M}^2f \equiv \mathbb{M}(\mathbb{M}f)$ is well-defined. If f is nonnegative and bounded, then $\mathbb{M}^nf \equiv \mathbb{M}(\mathbb{M}^{n-1}f)$ is defined for every $n \geq 1$, with $\mathbb{M}^0f \equiv f$ by definition. Using operator \mathbb{M} , we can express $(\gamma_n^N)_{0 \leq n \leq N}$ in terms of the process Π as stated in the following lemma.

Lemma 8. *For every $N \geq 0$, and $0 \leq n \leq N$, we have $\gamma_n^N = -c \sum_{k=0}^{n-1} (1 - \Pi_k^{(0)}) - (\mathbb{M}^{N-n}h)(\Pi_n)$.*

The next lemma shows how the optimal stopping problems can be rewritten in terms of the operator \mathbb{M} . It also conveys the connection between the truncated optimal stopping problems and the initial state Π_0 of the Π process.

Lemma 9. *We have (a) $V_0^N = (\mathbb{M}^N h)(\Pi_0)$ for every $N \geq 0$, and (b) $V_0 = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\Pi_0)$.*

Observe that since $\Pi_0 \in \mathcal{F}_0 = \{\emptyset, \Omega\}$, we have $\mathbb{P}\{\Pi_0 = \pi\} = 1$ for some $\pi \in S^M$. On the other hand, for every $\pi \in S^M$ we can construct a probability space $(\Omega, \mathcal{F}, \mathbb{P}_\pi)$ hosting a Markov process Π with the same dynamics as in (4) and $\mathbb{P}_\pi\{\Pi_0 = \pi\} = 1$. Moreover, on such a probability space, the preceding results remain valid. So, let us denote by \mathbb{E}_π the expectation with respect to \mathbb{P}_π and rewrite (8) as

$$-V_n(\pi) \triangleq \sup_{\tau \in C_n} \mathbb{E}_\pi Y_\tau \quad \text{and} \quad -V_n^N(\pi) \triangleq \sup_{\tau \in C_n^N} \mathbb{E}_\pi Y_\tau$$

for every $\pi \in S^M$. Then Lemma 9 implies that

$$V_0^N(\pi) = (\mathbb{M}^N h)(\pi) \quad \text{and} \quad V_0(\pi) = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\pi) \quad (9)$$

for every $\pi \in S^M$. Taking limits as $N \rightarrow \infty$ of both sides in $(\mathbb{M}^{N+1}h)(\pi) = \mathbb{M}(\mathbb{M}^N h)(\pi)$ and applying the monotone convergence theorem on the right-hand side yields $V_0(\pi) = (\mathbb{M}V_0)(\pi)$. Hence, we have shown the following result.

Proposition 10 (Optimality equation). *For every $\pi \in S^M$,*

$$V_0(\pi) = (\mathbb{M}V_0)(\pi) \equiv \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)\}. \quad (10)$$

Remark 11. *By solving $V_0(\pi)$ for any initial state $\pi \in S^M$, we capture the solution to the original problem since property (c) of Proposition 2 and (7) imply that $R^* = V_0(1 - p_0, p_0\nu_1, \dots, p_0\nu_M)$.*

B. Some properties of the value function.

Now, we reveal some important properties of the value function $V_0(\cdot)$ of (9). These results help us to establish an optimal solution for $V_0(\cdot)$, and hence an optimal solution for R^* , in the next subsection.

Lemma 12. *If $g : S^M \mapsto \mathbb{R}$ is a bounded concave function, then so is $\mathbb{T}g$.*

Proposition 13. *The mappings $\pi \mapsto V_0^N(\pi)$, $N \geq 0$ and $\pi \mapsto V_0(\pi)$ are concave.*

Proposition 14. *For every $N \geq 1$ and $\pi \in S^M$, we have*

$$V_0(\pi) \leq V_0^N(\pi) \leq V_0(\pi) + \left(\frac{\|h\|^2}{c} + \frac{\|h\|}{p} \right) \frac{1}{N}.$$

Since $\|h\| \triangleq \sup_{\pi \in S^M} |h(\pi)| < \infty$, $\lim_{N \rightarrow \infty} \downarrow V_0^N(\pi) = V_0(\pi)$ uniformly in $\pi \in S^M$.

Proposition 15. *For every $N \geq 0$, the function $V_0^N : S^M \mapsto \mathbb{R}_+$ is continuous.*

Corollary 16. *The function $V_0 : S^M \mapsto \mathbb{R}_+$ is continuous.*

Note that S^M is a compact subset of \mathbb{R}^{M+1} , so while continuity of $V_0(\cdot)$ on the interior of S^M follows from the concavity of $V_0(\cdot)$ by Proposition 12, Corollary 16 establishes continuity on all of S^M , including its boundary.

C. An optimal sequential decision strategy.

Finally, we describe the optimal stopping region in S^M implied by the value function $V_0(\cdot)$, and we present an optimal sequential decision strategy for our problem. Let us define for every $N \geq 0$,

$$\begin{aligned} \Gamma_N &\triangleq \{\pi \in S^M \mid V_0^N(\pi) = h(\pi)\}, \\ \Gamma_N^{(j)} &\triangleq \Gamma_N \cap \{\pi \in S^M \mid h(\pi) = h_j(\pi)\}, \quad j \in \mathcal{M}, \\ \Gamma &\triangleq \{\pi \in S^M \mid V_0(\pi) = h(\pi)\}, \\ \Gamma^{(j)} &\triangleq \Gamma \cap \{\pi \in S^M \mid h(\pi) = h_j(\pi)\}, \quad j \in \mathcal{M}. \end{aligned}$$

For each $j \in \{0\} \cup \mathcal{M}$, let $e_j \in S^M$ denote the unit vector consisting of zero in every component except for the j th component, which is equal to one. Note that e_0, \dots, e_M are the extreme points of the closed convex set S^M , and any vector $\pi = (\pi_0, \dots, \pi_M) \in S^M$ can be expressed in terms of e_0, \dots, e_M as $\pi = \sum_{j=0}^M \pi_j e_j$.

Theorem 17. *For every $j \in \mathcal{M}$, $(\Gamma_N^{(j)})_{N \geq 0}$ is a decreasing sequence of non-empty, closed, convex subsets of S^M . Moreover,*

$$\begin{aligned} \Gamma_0^{(j)} &\supseteq \Gamma_1^{(j)} \supseteq \dots \supseteq \Gamma^{(j)}, \\ \Gamma^{(j)} &\supseteq \{\pi \in S^M \mid h_j(\pi) \leq \min\{h(\pi), c(1 - \pi_0)\}\} \ni e_j, \\ \Gamma = \bigcap_{N=1}^{\infty} \Gamma_N &= \bigcup_{j=1}^M \Gamma^{(j)}, \quad \text{and} \quad \Gamma^{(j)} = \bigcap_{N=1}^{\infty} \Gamma_N^{(j)}, \quad j \in \mathcal{M}. \end{aligned}$$

Furthermore, $S^M = \Gamma_0 \supseteq \Gamma_1 \supseteq \dots \supseteq \Gamma \supsetneqq \{e_1, \dots, e_M\}$.

Lemma 18. *For every $n \geq 0$, we have $\gamma_n = -c \sum_{k=0}^{n-1} (1 - \Pi_k^{(0)}) - V_0(\Pi_n)$.*

Theorem 19. *Let $\sigma \triangleq \inf\{n \geq 0 \mid \Pi_n \in \Gamma\}$. (a) The stopped process $\{\gamma_{n \wedge \sigma}, \mathcal{F}_n; n \geq 0\}$ is a martingale.*

(b) *The random variable σ is an optimal stopping time for V_0 , and*

(c) $\mathbb{E} \sigma < \infty$.

Therefore, the pair (σ, d^*) is an optimal sequential decision strategy for (2), where the optimal stopping rule σ is given by Theorem 19, and, as in the proof of Lemma 4, the optimal terminal decision rule d^* is given by

$$d^* = j \quad \text{on the event } \{\sigma = n, \Pi_n \in \Gamma^{(j)}\} \quad \text{for every } n \geq 0.$$

Accordingly, the set Γ is called the *stopping region* implied by $V_0(\cdot)$, and Theorem 17 reveals its basic structure. We demonstrate the use of these results in the numerical examples of Sec. V.

Note that we can take a similar approach to prove that the stopping rules $\sigma_N \triangleq \inf\{n \geq 0 \mid \Pi_n \in \Gamma_{N-n}\}, N \geq 0$ are optimal for the truncated problems $V_0^N(\cdot), N \geq 0$ in (9). Thus, for each $N \geq 0$, the set Γ_N is called the stopping region for $V_0^N(\cdot)$: it is optimal to terminate the experiments in Γ_N if N stages are left before truncation.

V. SPECIAL CASES AND EXAMPLES

A. A. N. Shiryaev's sequential change detection problem.

Set $a_{0j} = 1$ for $j \in \mathcal{M}$ and $a_{ij} = 0$ for $i, j \in \mathcal{M}$, then the Bayes risk function (1) becomes $R(\delta) = \mathbb{P}\{\tau < \theta\} + c \mathbb{E}[(\tau - \theta)^+]$. This is the Bayes risk studied by Shiryaev [8, 9] to solve the sequential change detection problem.

B. Sequential multi-hypothesis testing.

Set $p_0 = 1$, then $\theta = 0$ a.s. and thus the Bayes risk function (1) becomes $R(\delta) = \mathbb{E}[c\tau + a_{\mu d} \mathbf{1}_{\{\tau < \infty\}}]$. This gives the sequential multi-hypothesis testing problem studied by Wald and Wolfowitz [10] and Arrow, Blackwell, and Girshick [11]; see also [12].

C. Two alternatives after the change.

In this subsection we consider the special case $M = 2$ in which we have only two possible change distributions, $f_1(\cdot)$ and $f_2(\cdot)$. We describe a graphical representation of the stopping and continuation regions for an arbitrary instance of the special case $M = 2$. Then we use this representation to illustrate geometrical properties of the optimal method (Sec. IV.C) via model instances for certain choices of the model parameters $p_0, p, \nu_1, \nu_2, f_0(\cdot), f_1(\cdot), f_2(\cdot), a_{01}, a_{02}, a_{12}, a_{21}$, and c .

Let the linear mapping $L : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be defined by $L(\pi_0, \pi_1, \pi_2) \triangleq (\frac{2}{\sqrt{3}}\pi_1 + \frac{1}{\sqrt{3}}\pi_2, \pi_2)$. Since $\pi_0 = 1 - \pi_1 - \pi_2$ for every $\pi = (\pi_0, \pi_1, \pi_2) \in S^2 \subset \mathbb{R}^3$, we can recover the preimage π of any point $L(\pi) \in L(S^2) \subset \mathbb{R}^2$. For every point $\pi = (\pi_0, \pi_1, \pi_2) \in S^2$, the coordinate π_i is given by the Euclidean distance from the image point $L(\pi)$ to the edge of the image triangle $L(S^2)$ that is *opposite* the image point $L(e_i)$, for each $i = 0, 1, 2$. For example, the distance from the image point $L(\pi)$ to the edge of the image triangle opposite the lower-left-hand corner $L(1, 0, 0) = (0, 0)$ is the value of the preimage coordinate π_0 . See Fig. ??.

Therefore, we can work with the mappings $L(\Gamma)$ and $L(S^2 \setminus \Gamma)$ of the stopping region Γ and the continuation region $S^2 \setminus \Gamma$, respectively. Accordingly, we depict the decision region for each instance in this subsection using the two-dimensional representation as in the right-hand-side of Fig. ?? and we drop the $L(\cdot)$ notation when labeling various parts of each figure to emphasize their source in S^2 .

Each of the examples in this section have the following model parameters in common:

$$p_0 = \frac{1}{50}, \quad p = \frac{1}{20}, \quad \nu_1 = \nu_2 = \nu_3 = \frac{1}{2}, \\ f_0 = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), f_1 = \left(\frac{4}{10}, \frac{3}{10}, \frac{2}{10}, \frac{1}{10}\right), f_2 = \left(\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}\right).$$

We vary the delay cost and false alarm/identification costs to illustrate certain geometrical properties of the continuation and stopping regions. See Figs. ??, ??, and ??.

These figures have certain features in common. On each subfigure there is a dashed line representing those states $\pi \in S^2$ at which $h_1(\pi) = h_2(\pi)$. Also, each subfigure shows a sample path of $(\Pi_n)_{n=0}^\sigma$ and the realizations of θ and μ for the sample. The shaded area, including its solid boundary, represents the optimal stopping region, while the unshaded area represents the continuation region.

Specifically, these figures show instances in which the $M = 2$ convex subsets comprising the optimal stopping region are connected (Fig. ??) and instances in which they are not (Figs. ?? and ??(a)). Fig. ??(b) shows an instance in which the continuation region is disconnected.

An implementation of the optimal strategy as described in Sec. IV.C is as follows: Initialize the statistic $\Pi = (\Pi_n)_{n \geq 0}$ by setting $\Pi_0 = (1 - p_0, p_0\nu_1, p_0\nu_2)$ as in part (c) of Proposition 2. Use the dynamics of (4) to update the statistic Π_n as each observation X_n is realized. Stop taking observations when the statistic Π_n enters the stopping region $\Gamma = \Gamma^{(1)} \cup \Gamma^{(2)}$ for the first time, possibly before the first observation is taken (i.e., $n = 0$). The optimal terminal decision is based upon whether the statistic Π_n is in $\Gamma^{(1)}$ or $\Gamma^{(2)}$ upon stopping. Each of the sample paths in Figs. ??, ??, and ?? were generated via this algorithm. As Fig. ?? shows, the sets $\Gamma^{(1)}$ and $\Gamma^{(2)}$ can intersect on their boundaries and so it is possible to stop in their intersection. In this case, either of the decisions $d = 1$ or $d = 2$ is optimal.

We use value iteration of the optimality equation (10) over a fine discretization of S^2 to compute $V_0(\cdot)$ and generate the decision region for each subfigure. The resulting discretized decision region is mapped into the plane via L . See [13, Ch. 3] for techniques of computing the value function via the optimality equation such as value iteration.

D. Three alternatives after the change.

In this subsection we consider the special case $M = 3$ in which we have three possible change distributions, $f_1(\cdot)$, $f_2(\cdot)$, and $f_3(\cdot)$. Here, the continuation and stopping regions are subsets of $S^3 \subset \mathbb{R}^4$. Similar to the two-alternatives case, we introduce the mapping of $S^3 \subset \mathbb{R}^4$ into \mathbb{R}^3 via $(\pi_0, \pi_1, \pi_2, \pi_3) \mapsto$

$$\left(\sqrt{\frac{3}{2}}\pi_1 + \frac{1}{2}\sqrt{\frac{3}{2}}\pi_2 + \frac{1}{2}\sqrt{\frac{3}{2}}\pi_3, \frac{3}{2}\sqrt{\frac{1}{2}}\pi_2 + \frac{1}{2}\sqrt{\frac{1}{2}}\pi_3, \pi_3\right).$$

Then we use this representation—actually a rotation of it—to illustrate in Fig. ?? an instance with the following model parameters:

$$p_0 = \frac{1}{50}, \quad p = \frac{1}{20}, \quad \nu_1 = \nu_2 = \nu_3 = \frac{1}{3}, \\ f_0 = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), \quad f_1 = \left(\frac{4}{10}, \frac{3}{10}, \frac{2}{10}, \frac{1}{10}\right), \\ f_2 = \left(\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}\right), \quad f_3 = \left(\frac{3}{10}, \frac{2}{10}, \frac{2}{10}, \frac{3}{10}\right), \\ c = 1, \quad a_{0j} = 40, \quad a_{ij} = 20, \quad i, j = 1, 2, 3.$$

Fig. ?? can be interpreted in a manner similar to the figures of the previous subsection. In this case, for every point $\pi = (\pi_0, \pi_1, \pi_2, \pi_3) \in S^3$, the coordinate π_i is given by the (Euclidean) distance from the image point $L(\pi)$ to the face of

the image tetrahedron $L(S^3)$ that is opposite the image corner $L(e_i)$, for each $i = 0, 1, 2, 3$.

REFERENCES

- [1] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall Inc., Englewood Cliffs, NJ, 1993.
- [2] T. L. Lai. Sequential analysis: some classical problems and new challenges. *Statist. Sinica*, 11(2):303–408, 2001.
- [3] I. V. Nikiforov. A generalized change detection problem. *IEEE Trans. Inform. Theory*, 41(1):171–187, 1995.
- [4] T. L. Lai. Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems. *IEEE Trans. Inform. Theory*, 46(2):595–608, 2000.
- [5] S. Dayanik, C. Goulding, and H. V. Poor. Bayesian sequential change diagnosis. <http://www.princeton.edu/~sdayanik/diagnosis.pdf>.
- [6] S. A. Klugman, H. H. Panjer, and G. E. Willmot. *Loss Models*. John Wiley & Sons Inc., New York, 1998.
- [7] Y. S. Chow, H. Robbins, and D. Siegmund. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Co., Boston, Mass., 1971.
- [8] A. N. Shiryaev. Optimal methods in quickest detection problems. *Teor. Verojatnost. i Primenen.*, 8:26–51, 1963.
- [9] A. N. Shiryaev. *Optimal Stopping Rules*. Springer-Verlag, New York, 1978.
- [10] A. Wald and J. Wolfowitz. Bayes solutions of sequential decision problems. *Ann. Math. Statistics*, 21:82–99, 1950.
- [11] K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17:213–244, 1949.
- [12] D. Blackwell and M. A. Girshick. *Theory of Games and Statistical Decisions*. Dover Publications Inc., New York, 1979. Reprint of the 1954 edition.
- [13] D. P. Bertsekas. *Dynamic Programming and Optimal Control. Vol. II*. Athena Scientific, Belmont, MA, second edition, 2001.